arXiv:1407.5609v1 [cs.DS] 21 Jul 2014

Efficient Algorithms for the Closest Pair Problem and Applications

Sanguthevar Rajasekaran, Sudipta Pathak
Department of CSE, University of Connecticut, Storrs, CT

July 2014

Abstract

The closest pair problem (CPP) is one of the well studied and fundamental problems in computing. Given a set of points in a metric space, the problem is to identify the pair of closest points. Another closely related problem is the fixed radius nearest neighbors problem (FRNNP). Given a set of points and a radius R, the problem is, for every input point p, to identify all the other input points that are within a distance of R from p. A naive deterministic algorithm can solve these problems in quadratic time. CPP as well as FRNNP play a vital role in computational biology, computational finance, share market analysis, weather prediction, entomology, electro cardiograph, N-body simulations, molecular simulations, etc. As a result, any improvements made in solving CPP and FRNNP will have immediate implications for the solution of numerous problems in these domains. We live in an era of big data and processing these data take large amounts of time. Speeding up data processing algorithms is thus much more essential now than ever before. In this paper we present algorithms for CPP and FRNNP that improve (in theory and/or practice) the best-known algorithms reported in the literature for CPP and FRNNP. These algorithms also improve the best-known algorithms for related applications including time series motif mining and the two locus problem in Genome Wide Association Studies (GWAS).

1 Introduction

The closest pair problem (CPP) has a rich history and has been extensively studied. On an input set of n points, the problem is to identify the closest pair of points. A straight forward algorithm for CPP takes quadratic (in n) time. Most of the algorithms proposed in the literature are concerned with the Euclidean space. In the rest of this paper, unless otherwise mentioned, we imply the Euclidean space. In his seminal paper, Rabin proposed a randomized algorithm with an expected run time of $O(n)$ [20] (where the expectation is in the space of all possible outcomes of coin flips made in the algorithm). Rabin's algorithm used the floor function as a basic operation. In 1979, Fortune and Hopcroft

presented a deterministic algorithm with a run time of $O(n \log\log n)$ assuming that the floor operation takes $O(1)$ time [10]. Both of these algorithms assume a constant-dimensional space (and the run times have an exponential dependency on the dimension). The algorithm of Preparata and Shamos for points in 2D is deterministic and runs in $O(n \log n)$ time [19]. Yao has proven a lower bound of $\Omega(n \log n)$ on the algebraic decision tree model (for any dimension) [26]. This lower bound holds under the assumption that the floor function is not allowed. The algorithm of Khuller and Matias is also randomized and has an expected linear run time utilizing the floor operation [12]. There are two steps in the algorithm. In the first step, the distance between the closest pair of points is estimated within a factor of 3. In the second step, the neighborhood of each point p is explored to identify those points that are within a distance of e from p, where e is the estimate that step 1 comes up with. Using this neighborhood information, the correct pair is identified.

One of the major issues with the above algorithms is the fact that their run times are exponentially dependent on the dimension. For example, the expected run time of [12]'s algorithm is $O(3^d n)$ on n points from a d-dimensional space. So, even for a moderate value of d, the algorithm may be very slow in practice. There are numerous applications for which the dimension is very large. In the following sections, we consider two such application domains.

Time series motif mining (TSMM) is a crucial problem that can be thought of as CPP in a large dimensional space. In one version of the TSMM problem, we are given a sequence S of real numbers and an integer ℓ. The goal is to identify two subsequences of S of length ℓ each that are the most similar to each other (from among all pairs of subsequences of length ℓ each). These most similar subsequences are referred to as *time series motifs*. Let C be a collection of all the ℓ-mers of S. (An ℓ-mer is nothing but a contiguous subsequence of S of length ℓ). Clearly, the ℓ-mers in C can be thought of as points in $\Re^\ell$. As a result, the TSMM problem is the same as CPP in $\Re^\ell$. Any of the above mentioned algorithms can thus be used to solve the TSMM problem. A typical value for ℓ of practical interest is several hundreds (or more). For these values of ℓ, the above algorithms ([20],[10],[19],[12]) will take an unacceptable amount of time (because of the exponential dependence on the dimension). Designing an efficient practical and exact algorithm for the TSMM problem remains an ongoing challenge.

Mueen, et al. have presented an elegant exact algorithm called MK for TSMM [15]. MK improves the performance of the brute-force algorithm with a novel application of the triangular inequality. MK is currently the best-performing algorithm in practice for TSMM. A number of probabilistic as well as approximate algorithms are also known for solving this problem (see e.g., [3, 5, 11, 13, 14, 21, 22]). For instance, the algorithm of [5] exploits algorithms proposed for finding (ℓ, d)-motifs from biological data. The idea here is to partition the time series data into frames of certain width. Followed by this, the mean value in each frame is computed. This mean is quantized into four intervals and as a result, the original time series data is converted into a string of characters from an alphabet of size 4. Finally, any (ℓ, d)-motif finding algorithm

is applied on the transformed string to identify the time series motifs.

An application of great interest in bioinformatics is Genome Wide Association Study (GWAS). A lot of effort has been spent to identify mappings between phenotypical traits and genomic data. Due to the advent of next generation high throughput sequencing technologies, nowadays it is possible to study the genomic structure of individuals in detail. Single Nucleotide Polymorphisms (SNPs) are positions in the genome where nucleotides vary among individuals [7]. SNPs in the human genome are considered to be responsible for different phenotypical traits. In GWAS, two different problems have been focused on. In single locus association study, researchers try to find out the association between phenotypical traits and individual SNPs. In two locus association study, the goal is to figure out the association between pairs of SNPs and phenotypical traits. A major task in this study is that of identifying the most correlated pair of SNPs. Two locus associations are also known as gene-gene interactions. Such interactions are believed to be major factors responsible for many complex phenotypical traits [16, 2, 6, 8, 9, 25].

Given that the number of SNPs found in humans is 10^5 to 10^7, a brute force way of scanning through every possible pair of SNPs to identify the most correlated pair is not feasible in practice. A number of algorithms for the two locus problem can be found in the literature. For instance, genetic algorithms are used in [17] and [24]. The algorithms proposed in [27] and [28] take $O(n^2m + nm^2)$ time, where n is the number of SNPs and m is the number of subjects. An algorithm with an expected run time of $O(mn^{1+\epsilon})$, where $0 < \epsilon < 1$ is a constant, has been presented in [1]. This algorithm exploits an algorithm known for the Light Bulb Problem [18] and Locality Sensitive Hashing (LSH) [4].

In this paper we present efficient algorithms for CPP. Our algorithms improve the results reported in several papers including [15], [5], and [1]. Our main contributions can be summarized as follows:

1. We improve the CPP algorithm of [15] by introducing a novel idea. Specifically, [15]'s algorithm is based on two basic ideas. We contribute a third idea that results in an improvement of the run time for exact TSMM by a factor of around 1.5. We also show how to extend our algorithm to solve the fixed radius nearest neighbors problem (FRNNP).

2. We present an algorithm for CPP when the domain of interest has strings of characters (from a finite alphabet) and the metric is Hamming distance. It turns out that MK does not perform well for the case of Hamming distance, and to be fair, we note that the authors of MK do not claim it might. A comparison of our algorithm with MK reveals that our algorithm outperforms MK (by a factor of around 200). Our algorithm can also be used in approximate TSMM. Specifically, instead of using (ℓ, d)-motif search algorithms, our algorithm can be used in [5]. In this case, the run time of the algorithm in [5] will improve significantly, since exact algorithms for solving the (ℓ, d)-motif search problem take time that is exponential in ℓ and d. Our algorithm is a modified version of the light bulb algorithm of [18].

3. The light bulb algorithm of [18] finds the most correlated pair of bulbs. The light bulb problem can be thought of as CPP in the space of **binary** strings with Hamming distance as the metric. We extend this algorithm when the strings are from an arbitrary (finite) alphabet. More importantly, we present an algorithm for finding the **least** correlated pair of strings (from an arbitrary alphabet). The algorithm of [1] also solves this problem utilizing Locality Sensitive Hashing (LSH) [4]. Our algorithm does not use LSH. Instead, it uses a novel deterministic mapping function that we have come up with.

4. Using the above algorithm for finding the least correlated pair, we present a novel algorithm for solving the two locus GWAS problem. An experimental comparison reveals that our algorithm is four times faster than the algorithm of [1]. We note here that the authors of [1] use Pearson's correlation coefficient to measure the similarity between a pair of SNPs, whereas we use the complement of the Hamming distance as the measure of similarity.

The rest of this paper is organized as follows. In Section 2 we provide some notations. In Section 3 we present our improved algorithm for CPP (called MPR) and compare it with MK. In this section we also provide an analysis of MK and experimentally compare the performances of MK and MPR. Section 4 deals with the case of character strings and Hamming distance. Specifically, we show how to modify the light bulb algorithm of [18] to get an algorithm for finding the most correlated pair of strings from an arbitrary finite alphabet. We compare this algorithm with MK experimentally. In Section 5 we present an algorithm for finding the least correlated pair of strings. This algorithm is based on a novel mapping function that we have come up with. Section 6 is devoted to the problem of two locus association in GWAS. In particular, we present a novel algorithm for this problem and compare our algorithm with that of [1]. Some concluding remarks are given in Section 7.

2 Notations and Definitions

Let $T = a_1, a_2, \ldots, a_n$ be a sequence of real numbers (or characters from a finite alphabet). An ℓ-mer of T is nothing but a subsequence of T of ℓ contiguous elements of T. The ℓ-mers of T are $T_i = a_i, a_{i+1}, \ldots, a_{i+\ell-1}$, for $1 \leq i \leq (n - \ell + 1)$.

If the elements of T are real numbers, then the Euclidean distance between T_i and T_j, denoted as $d(T_i, T_j)$, is $\sqrt{\sum_{k=0}^{\ell-1}(a_{i+k} - a_{j+k})^2}$.

If the elements of T are characters from an alphabet Σ, then the Hamming distance between T_i and T_j, denoted as $d(T_i, T_j)$, is $\sum_{k=0}^{\ell-1} \delta(a_{i+k}, a_{j+k})$ where $\delta(a, b) = 1$ if $a \neq b$ and $\delta(a, b) = 0$ if $a = b$ (for any $a, b \in \Sigma$). A sequence of characters can be thought of as a string of characters, since we can obtain a string from the sequence by concatenating the characters. Thus we'll use the terms 'a sequence of characters' and 'a string of characters' interchangeably.

Let $A = a_1, a_2, \ldots, a_n$ and $B = b_1, b_2, \ldots, b_n$ be two sequences of characters. Also, let the Hamming distance between A and B be d. Then, by the *number of matches* between A and B we mean $n - d$. Also, the *correlation* between A and B is defined to be $\frac{n-d}{n}$.

3 Time Series Motif Mining Algorithm

The input for this problem are a sequence $T = a_1, a_2, \ldots, a_n$ and an integer ℓ. The goal is to find two ℓ-mers of T that are the closest to each other (from among all the pairs of ℓ-mers of T). A general version of this problem is one where the input consists of n points from $\Re^\ell$ and we want to identify the two closest points. A straight forward algorithm will compute the distance between every pair of ℓ-mers and output the pair with the least distance. Since we can compute the distance between two ℓ-mers in $O(\ell)$ time, this simple algorithm for TSMM will run in a total of $O(n^2\ell)$ time.

3.1 MK Algorithm

The MK algorithm of [15] speeds up the brute force method by pruning off a large number of pairs that cannot possibly be the closest. There are two main ideas used in MK. The first idea in the algorithm is to speedup the computation of distances. Let $x = x_1, x_2, \ldots, x_\ell$ and $y = y_1, y_2, \ldots, y_\ell$ be any two ℓ-mers. To compute the distance between x and y, the algorithm keeps adding $(x_i - y_i)^2$ for $i = 1, 2, \ldots, \ell$. When the sum exceeds δ^2, this pair is immediately dropped (without completing the rest of the distance computation). This technique is known as *early abandoning*.

The second idea in MK uses the triangular inequality in a novel way. Let x and y be any two ℓ-mers. At any stage in the algorithm, we have an upper bound δ on the distance between the closest pair of ℓ-mers. If $d(x, y)$ can be inferred to be greater than δ, then we can drop the pair (x, y) from future consideration (since this pair cannot be the closest). Ideally, we would like to calculate $d(x, y)$ exactly for every pair of ℓ-mers x and y. But this will take too much time. MK circumvents this problem by **estimating** the distance between x and y via the triangular inequality. In particular, a random reference ℓ-mer r is chosen and the distance between each ℓ-mer and r is computed. The ℓ-mers are kept in an ascending order of their distances to r. From thereon, $d(r, y) - d(r, x)$ is used as a lower bound on $d(x, y)$. If this lower bound is $> \delta$, then (x, y) is dropped from future consideration.

The above algorithm is generalized to employ multiple reference ℓ-mers. The use of multiple references speeds up the algorithm.

3.2 An Analysis of the MK Algorithm and Our New Idea

In this section we provide an (informal) analysis of the MK algorithm to explain why the algorithm has a very good performance. Specifically, if we choose

multiple random reference points, the algorithm achieves a much better run time than having a single reference point. We explain why this is the case.

For ease of understanding consider the 2D Euclidean space. The analysis can be extended to points in $\Re^\ell$. For any two ℓ-mers x and y, the closer $d(r,y) - d(r,x)$ is to $d(x,y)$, the better will be our estimate and hence the better will be our chance of dropping (x,y) (if (x,y) is not the closest pair). It turns out that the quality of the lower bound $d(r,y) - d(r,x)$ is decided by two factors: 1) the angle $\angle rxy$ and 2) $d(r,x)$. We illustrate this with an example. Let $x = (0,0)$ and $y = (1,0)$. Consider a reference point $r_1 = (1,1)$ (Figure 1(a)). Note that r_1 is at a distance of $\sqrt{2}$ from x. In this case $d(r_1,x) - d(r_1,y) = 0.414$. Also, $\angle r_1xy = 45°$. Let r_2 be the point we get by keeping the distance between the reference point and x the same, but changing this angle to 30° (Figure 1(b)). In this case, $d(r_2,x) - d(r_2,y)$ improves to 0.6722. As another example, if the reference point r lies on the perpendicular bisector of x and y, then $d(r,x) - d(r,y) = 0$.

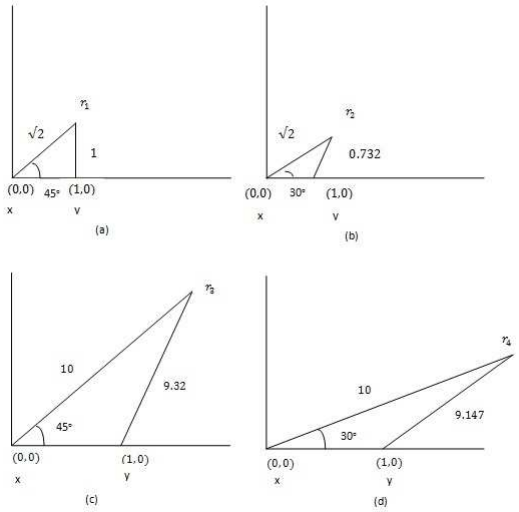

Figure 1: Reference Points

For any two input points x and y, if we pick multiple reference points randomly, then we would expect that at least one of these reference points r will be such that the angle $\angle rxy$ will be such that $d(r,x) - d(r,y)$ will be 'large'. In contrast, if we have only one reference point, for some pairs of points the corresponding angles may be 'good', but for a good percentage of the pairs, the angles may not be 'good'. (A reference point r is ''good' if $d(r,x) - d(r,y)$ is close to $d(x,y)$).

The effect of $d(r,x)$ on $d(r,x) - d(r,y)$ can be seen with the same examples. Consider the example of Figure 1(a). Assume that we keep the angle the same but increase $d(r_1,x)$ to 10 and get the reference point r_3 (Figure 1(c)). In this case, $d(r_3,x) - d(r_3,y)$ improves to 0.6803. Also, in the example of Figure 1(b),

say we keep the angle the same but increase $d(r_2, x)$ to 10 and get the reference point r_4. In this case, $d(r_4, x) - d(r_4, y)$ improves to 0.8526. Of course, if the reference point r lies on the perpendicular between x and y, then however large $d(r, x)$ could be, $d(r, x) - d(r, y)$ will continue to be zero! However, the probability of this happening is low. For a given angle θ, we can compute the limit of $d(r, x) - d(r, y)$ as $d(r, x)$ tends to ∞. For instance when the angle is 45° (Figure 1(a)), this limit is $\frac{1}{\sqrt{2}} \approx 0.707$.

3.3 Our algorithm

Our proposed new algorithm indeed exploits the relationship between $d(r, x)$ and $d(r, x) - d(r, y)$. In particular, we pick a collection C of random reference points and project each of these points out by multiplying each coordinate value of each point by a factor of f. For example, f could be 10. The rest of the algorithm is the same as MK.

A pseudocode for our algorithm, called *Motif discovery with Projected Reference points (MPR)*, is given below.

Algorithm MPR

Input: $T = a_1, a_2, \ldots, a_n$ and an integer ℓ, where each a_i is a real number (for $1 \leq i \leq n$). Input are also q and f, where q is the number of references and f is the projection factor.

Output: The two closest ℓ-mers of T.

1) Pick q random ℓ-mers of T as references; Project these references by multiplying each element in each ℓ-mer by f. Let these projected references be $r_1, r_2, \ldots, r_q$.

2) Compute the distance between every ℓ-mer of T and every projected reference ℓ-mer.

3) Sort the ℓ-mers of T with respect to their distances to r_1. Let the sorted ℓ-mers be $p_1, p_2, \ldots, p_{n-\ell+1}$.

4) Let $\delta = \infty$; Let $answer = (0, 0)$;

5) **for** $i := 1$ **to** $(n - \ell + 1)$ **do**
    for $j := (i + 1)$ **to** $(n - \ell + 1)$ **do**
        $failure := false$;
        for $k := 1$ **to** q **do**
            if $d(r_k, p_j) - d(r_k, p_i) > \delta$ **then**
                $failure := true$; **exit**;
        if $failure$ **then exit else**
            Compute $d(p_i, p_j)$;
            if $d(p_i, p_j) < \delta$ **then**
                $\delta := d(p_i, p_j)$; $answer = (i, j)$;

6) Output (i, j).

Observation: Please note that even though the above algorithm has been presented for solving the TSMM problem, it is straight forward to extend it to solve CPP in $\Re^\ell$.

3.4 An analysis of our algorithm

In this section we show why our idea of projecting reference points improves the performance of the algorithm. Let the input points be from $\Re^d$ for some integer d. Consider any two input points A and B. Let R be any reference point. Note that any three points are coplanar. Consider any hyperplane $\mathcal{H}$ containing A, B, and R. If we multiply every coordinate of R by the same number, then the resultant point will also lie in $\mathcal{H}$. This is because the equation defining $\mathcal{H}$ will be of the form $a_1x_1 + a_2x_2 + \cdots + a_dx_d = 0$. Thus, in order to see how $d(R, A) - d(R, B)$ changes with a scaling of R, it suffices to consider the case that these three points are in 2D.

Without loss of generality let A be $(0, 0)$ and B be $(c, 0)$, for some real number c. There are two cases to consider for the position of R relative to A and B: 1) R is to the right of the perpendicular bisector of A and B; 2) R is to the left of the perpendicular bisector between A and B. These two cases are illustrated in Figures 2(a) and 2(b), respectively. Note that when R lies on the perpendicular bisector of A and B, $d(R, A) - d(R, B)$ will be zero.

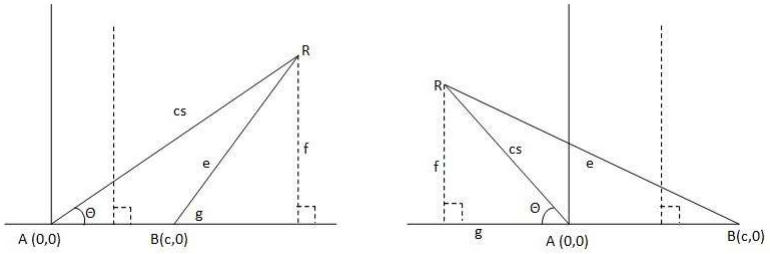

Figure 2: The effect of scaling on reference points

In case 1, let $d(A, R)$ be cs, s being a scaling factor. $f = cs\sin\theta$, $g = cs\cos\theta - c$, and $e = \sqrt{f^2 + g^2}$. As a result, $e = \sqrt{c^2s^2 + c^2 - 2c^2s\cos\theta} = cs\sqrt{1 + \frac{1}{s^2} - \frac{2\cos\theta}{s}}$. Using the fact that $(1 - u)^n \approx 1 - nu$ (when $nu << 1$), $e \approx cs + \frac{c}{2s} - c\cos\theta$. Thus, $d(A, R) - d(R, B) \approx c\cos\theta - \frac{c}{2s}$. Clearly, when c and θ are the same, the value of $d(R, A) - d(R, B)$ increases when s increases.

In case 2, let $d(A, R)$ be cs, for a scaling factor of s. Clearly, $f = cs\sin\theta$, $g = cs\cos\theta$. Thus, $e = \sqrt{f^2 + (g + c)^2}$. Also, $e = \sqrt{c^2s^2 + c^2 + 2c^2s\cos\theta} = cs\sqrt{1 + \frac{1}{s^2} + \frac{2\cos\theta}{s}}$. Using the approximation mentioned in case 1, we see that $e \approx cs + \frac{c}{2s} + c\cos\theta$. Therefore, $d(R, B) - d(R, A) \approx \frac{c}{2s} + c\cos\theta$. In this case, when c and θ are the same, the value of $d(R, B) - d(R, A)$ increases when s decreases.

But for a given reference point, and two input points A and B, we do not know which of the two cases will hold. But we can expect that half of the randomly chosen reference points will fall under case 1 and the other half will be expected to fall under case 2. If we only employ a scaling factor s that is

greater than one, then, for an expected half of the reference points we expect to see an improvement in the estimate of a lower bound for $d(A, B)$. This explains why our algorithm performs better than MK.

The above analysis can also be used to better understand the MK algorithm.

3.5 Fixed radius nearest neighbors problem (FRNNP)

In this problem we are given n points $a_1, a_2, \ldots, a_n$ in $\Re^\ell$ and a radius R (which is a real number) and the problem is to identify the R-neighborhood of each input point. If p is an input point, its R-neighborhood is defined to be the set of all input points that are within a distance of R from p. FRNNP has numerous applications. One of the applications of vital importance is that of molecular simulations.

We can modify MPR to solve this problem as well. The modified version is given below. Let $N(i)$ denote the R-neighborhood of $a_i, 1 \leq i \leq n$.

Algorithm MPR-FRNNS

Input: $T = a_1, a_2, \ldots, a_n$ and R, where each a_i is a point in $\Re^\ell$ (for $1 \leq i \leq n$) and R is a real number. Input are also q and f, where q is the number of references and f is the projection factor.

Output: $N(i)$, for $1 \leq i \leq n$.

1) Pick q random points of T as references; Project these references by multiplying each coordinate of each reference point by f. Let these projected references be $r_1, r_2, \ldots, r_q$.
2) Compute the distance between every point of T and every projected reference point.
3) Sort the points of T with respect to their distances to r_1. Let the sorted points be $p_1, p_2, \ldots, p_n$.
4) for $i := 1$ **to** n **do** $N(i) := \emptyset$;
5) for $i := 1$ **to** n **do**
    for $j := (i + 1)$ **to** n **do**
        $failure := false$;
        for $k := 1$ **to** q **do**
            if $d(r_k, p_j) - d(r_k, p_i) > R$ **then**
                $failure := true$; **exit**;
        if $failure$ **then exit else**
            Compute $d(p_i, p_j)$;
            if $d(p_i, p_j) \leq R$ **then** add j to $N(i)$ and
            add i to $N(j)$;
6) Output $N(i)$, for $1 \leq i \leq n$.

3.6 An experimental comparison of MK and MPR

A typical algorithm in the literature for CPP has two phases. In the first phase pairs of points that cannot possibly be the closest are eliminated. In the second phase distance is computed between every pair of points that survive the first phase. The time spent in the first phase is typically very small and

hence is negligible (compared to the time spent in the second phase). Also, the time needed to process the pairs in the second phase is linear in the number of surviving pairs. As a result, it suffices to report the number of pairs (to be processed in the second phase) as a measure of performance (see e.g., [1]). In this paper also we use this measure of performance throughout.

We have experimentally compared the performance of MK and MPR on different data sets. The machine we have used has an Intel(R) Core(TM) i7-2640M 2.8 GHz CPU with 8GB RAM running Windows 7 (64 bit). The same machine has been used for all the experiments reported in this paper.

As mentioned in [15], random walk data set is the most difficult case for time series mining algorithms since the probability of the existence of very close motif pairs is very low. We have also used the same data for our comparison. In particular, we have used 10 different random walk data sets of sizes ranging from 10K to 100K. We have also varied the motif length to see how the performances change. Our algorithm performs better than MK for higher motif lengths. Both the algorithms have been run 10 times and the averages computed. We do not perform any comparison with the brute force method as that has already been done in [15].

Dataset Size	Pairs (MK)	Runtime in sec(MK)	Pairs (MPR)	Runtime in sec(MPR)
$10,000$	1.0078×10^7	18.5	0.7634×10^7	10.1
$20,000$	1.1858×10^7	60.7	0.7650×10^7	40.8
$30,000$	2.0001×10^7	60.9	1.3793×10^7	42.3
$40,000$	3.1785×10^7	62.3	2.1476×10^7	43.5
$50,000$	4.8031×10^7	101.5	3.2897×10^7	51.8
$60,000$	5.1409×10^7	102.1	3.5454×10^7	52.3
$70,000$	1.4083×10^8	288.9	0.9388×10^8	167.8
$80,000$	1.7896×10^8	340.9	1.1930×10^8	232.8
$90,000$	2.1198×10^8	550.3	1.4519×10^8	341.0
$100,000$	3.1949×10^8	772.5	2.1587×10^8	520.9

Table 1: Number of Pairs and Runtime comparison: Euclidean case

In Table 1 we show the number of pairs processed (in the second phase) in MK and MPR. The size (i.e., the length) of the time series data varies from 10K to 100K, the motif length being 1024. The following parameter values have been used: $q = 10$ and $f = 10$. From this table we see that MK processes around 1.5 times the number of pairs processed by MPR. Figure 3 presents a graphical plot of the runtime requirements of MK and MPR algorithms. This figure shows that the run time of MK is around 1.5 times the run time of MPR.

This improvement is quite significant for the following reason: There are two ideas used in MK, namely, early abandoning and the use of random reference points. As the authors point out in [15], the difference between using early abandoning alone and both the ideas is small, especially on random walk data sets. However, MK performs much better than using early abandoning alone on real datasets.

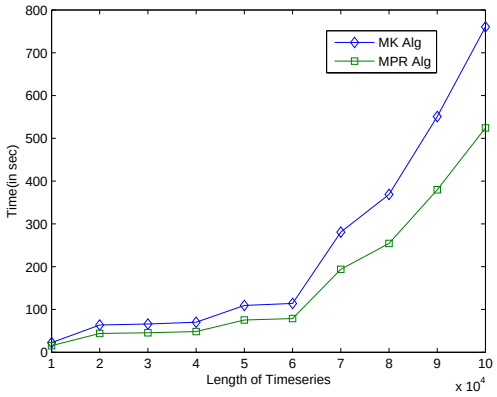

Figure 3: Runtime Comparison between MK and MPR: Euclidean case

Table 2 and Figure 4 show how the number of pairs reduces with an increase in the number q of the reference points for MK and MPR algorithms. From these, we note that as q increases, the difference between MK and MPR widens.

The bar graph in Figure 4 pictorially represents this comparison. The blue and red bars represent the number of pairs processed by MK and MPR, respectively.

We have run both MK and MPR algorithms on some of the real data sets from `http://www.cs.ucr.edu/~mueen/MK/`. Table 3 shows the performance of MK and MPR algorithms. On dataset 1, MPR is around 2 times faster than MK and on dataset 2, MPR is around 1.5 times faster than MK. In this Table, '-' indicates that the algorithm did not stop within 40 minutes.

4 The Case of Character Strings

In this section we consider the CPP when the space is one of character strings and the metric is Hamming distance. An algorithm for this version of CPP has numerous applications including approximate TSMM (see e.g., [5]). When the alphabet is $\{0, 1\}$, the light bulb algorithm of [18] can be used to solve this problem. In this section we show how to modify the light bulb algorithm for the case of generic alphabets. Before presenting details on the modification, we provide a brief summary of the light bulb problem.

No of Ref	No of Pairs(MK)	Runtime(MK)	No of Pairs(MPR)	Runtime(MPR)
1	1.7157462×10^7	18.93 sec	1.2112471×10^7	13.17 sec
2	0.7806227×10^7	10.43 sec	0.3773235×10^7	5.03 sec
3	0.3500211×10^7	5.74 sec	0.1534813×10^7	3.15 sec
4	0.1769827×10^7	4.29 sec	0.0896315×10^7	2.23 sec
5	0.1255878×10^7	3.15 sec	0.0458978×10^7	2.16 sec
6	0.0885219×10^7	2.88 sec	0.0314989×10^7	1.90 sec
7	0.0725813×10^7	2.73 sec	0.0273630×10^7	1.73 sec
8	0.0547364×10^7	2.69 sec	0.0216532×10^7	1.68 sec
9	0.0460080×10^7	2.45 sec	0.0179179×10^7	1.60 sec
10	0.0410578×10^7	2.46 sec	0.0178203×10^7	1.51 sec
20	0.0297705×10^7	2.37 sec	0.0173253×10^7	1.32 sec
30	0.0353774×10^7	2.06 sec	0.0310319×10^7	1.18 sec
40	0.0432340×10^7	1.99 sec	0.0405625×10^7	1.14 sec
50	0.0523544×10^7	1.96 sec	0.0502900×10^7	1.09 sec
60	0.0615568×10^7	2.00 sec	0.0598589×10^7	1.12 sec
70	0.0711500×10^7	2.01 sec	0.0696858×10^7	1.12 sec
80	0.0806084×10^7	2.12 sec	0.0794800×10^7	1.12 sec
90	0.0903655×10^7	2.17 sec	0.0892789×10^7	1.19 sec
100	0.1001488×10^7	2.19 sec	0.0991276×10^7	1.19 sec

Table 2: Number of pairs as a function of the number of references: Euclidean case

4.1 The light bulb problem

The light bulb problem is that of identifying the most correlated pair of bulbs from out of n given bulbs
$b_1, b_2, \ldots, b_n$. This problem is solved by observing the state of each bulb in t discrete time steps (for some relevant value of t). The states of bulb i in these t time steps can be represented as a vector $\hat{b}_i = (b_1^i, b_2^i, \ldots, b_t^i)$ (for $1 \leq i \leq n$). We can think of $\hat{b}_i$ as a sample from a probability distribution that governs the state of bulb i. It can be shown that if t is sufficiently large (e.g., $\Omega(\log n)$), then the pair of bulbs that is the most correlated in the samples is also the most correlated pair with high probability. Thus the light bulb problem can be stated as follows: We are given n Boolean vectors $\hat{b}_1, \hat{b}_2, \ldots, \hat{b}_n$. The problem is to find the pair of vectors that are the most similar (i.e., the Hamming distance

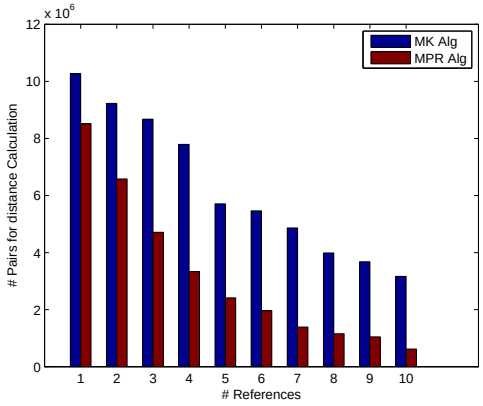

Figure 4: Number of pairs as a function of the number of reference points: Euclidean case

Dataset	Length	No. of Pairs (MK)	Runtime(MK)	No. of Pairs (MPR)	Runtime(MPR)
1	33,021	1.2719×10^7	43.944 sec	0.6112×10^7	25.214 sec
2	18,667	2.1927×10^8	485.591 sec	1.4570×10^8	355.531 sec
3	78,254	19.0220×10^8	6018.072 sec	1.5027×10^9	3857.739 sec

Table 3: Comparison with Real Data sets.
Data set 1: $RepeatedInsectBehaviorDataset1$;
Data set 2: $RepeatedInsectBehaviorDataset2$;
Data set 3: $RepeatedInsectBehaviorDataset3$

between them is the smallest). From hereon, we will use this formulation of the problem.

Note that, given two vectors, we can find the Hamming distance between them in $O(t)$ time. A straight forward algorithm to identify the most correlated pair of bulbs takes $O(n^2 t)$ time. This algorithm computes the Hamming distance between every pair of bulbs. The algorithm of [18] takes subquadratic time. In particular, the expected run time of this algorithm is $O\left(n^{1+\frac{\log p_1}{\log p_2}} \log^2 n\right)$ assuming that $t = O(\log n)$. Here, p_1 is the correlation between the most correlated pair of bulbs and p_2 is the correlation between the second most correlated pair of bulbs. Note that if the correlation between two bulbs i and j is p_{ij} then the expected Hamming distance between $\hat{b}_i$ and $\hat{b}_j$ is $t(1 - p_{ij})$. Equivalently, the similarity (i.e., the number of matches) between $\hat{b}_i$ and $\hat{b}_j$ is tp_{ij}.

4.2 The light bulb algorithm

Consider a matrix M of size $n \times t$, such that the ith row of M is $\hat{b}_i$, for $1 \leq i \leq n$. The algorithm of [18] iteratively collects pairs of bulbs that are candidates to be the most correlated. Once it collects enough pairs, it computes the distance between each pair in this collection and outputs the closest. There are $O\left(n^{\frac{\log p_1}{\log p_2}} \log n\right)$ iterations in the algorithm and in each iteration, some candidate pairs are generated and added to the collection C. In any iteration, the algorithm picks $c \log n$ columns of M at random (for some constant c). The rows are sorted based on the characters in the randomly chosen columns. As a result of this sorting, the bulbs get partitioned into buckets such that all the bulbs with equal values (in the $c \log n$ random columns) fall into the same bucket. A pair of bulbs (a, b) will get added to C in any iteration if they fall into the same bucket in this iteration.

The authors of [18] show that after $O\left(n^{\frac{\log p_1}{\log p_2}} \log n\right)$ iterations, C will have the most correlated pair of bulbs with high probability (i.e., with a probability of $1 - n^{-\Omega(1)}$). The above algorithm has been proposed for the case of binary strings. We can modify this algorithm to handle the case of an arbitrary (finite) alphabet and get the following theorem.

Theorem 4.1. *Let M be a matrix of size $n \times t$. Each entry in this matrix is an element from some set Σ of cardinality σ. We can find the most correlated pair of columns of M in an expected $O\left(n^{1+\frac{\log p_1}{\log p_2}} \frac{\log^2 n \log \sigma}{w}\right)$ time where p_1 is the correlation between the most correlated pair of columns, p_2 is the correlation between the second most correlated pair of columns, and w is the word length of the machine. This expected run time will be $O\left(n^{1+\frac{\log p_1}{\log p_2}} \log^2 n\right)$ if we use a general sorting algorithm. (Here correlation is based on Hamming distance. For example, p_1 is the largest fraction of rows in which any two columns agree).*

Proof: The only difference in the algorithm is that instead of sorting binary strings we will have to sort strings from an arbitrary alphabet. Without loss of generality, let $\Sigma = \{0, 1, 2, \ldots, \sigma - 1\}$ be the alphabet under concern. In the original algorithm, one has to sort n $(c \log n)$-bit integers in every iteration. For a generic alphabet, we have to sort n $(c \log n \log \sigma)$-bit integers. If one uses any comparison based sorting algorithm, this sorting takes $O(n \log n)$ time. If we use an integer sorting algorithm, this sorting can be done in $O\left(\frac{cn \log n \log \sigma}{w}\right)$ time where w is the word length of the machine. This is the time spent in each iteration of the algorithm. Therefore, the total expected run time is $O\left(n^{1+\frac{\log p_1}{\log p_2}} \frac{\log^2 n \log \sigma}{w}\right)$. □

Let this modified version of the light bulb algorithm be called MLBA.

4.3 An experimental comparison

The algorithm of [5] for approximate TSMM partitions the input time series data T based on a window of size w (for an appropriate value of w), computes

the mean of every window, and discretizes the mean into four possible values. As a result, the time series data is transformed into a string T' of characters form the alphabet $\{1, 2, 3, 4\}$. It then uses any (ℓ, d)-motif finding algorithm to find the motifs in T'. However, all the exact algorithms for finding (ℓ, d)-motifs take time that is exponential on ℓ and d. Note that the last step of finding (ℓ, d) motifs can be replaced with a problem of finding time series motifs in T' which is nothing but CPP in the domain of strings of characters, the motif length being ℓ.

One could employ MK to solve CPP in the domain of character strings. The only difference is that we have to replace Euclidean distance with Hamming distance. We have implemented this algorithm. It turns out that MK does not perform well for the case of Hamming distance. To be fair, the authors of MK have not tested MK for this case. We have compared MK with MLBA and the results are shown in Table 4. As this Table reveals, MLBA is around 200 times faster than MK. It is also clear that if we employ MLBA in place of (l, d)-motif finding algorithms, the performance of the approximate TSMM algorithm given in [5] will improve significantly.

Dataset Size	Pairs(MK)	Runtime in sec(MK)	Pairs(MLBA)	Runtime in sec(MLBA)
$1,000$	176,799	0.0067	2,012	0.0002
$2,000$	1,981,048	10.1	7,849	0.0003
$3,000$	5,783,132	16.5	17,795	0.0005
$4,000$	11,586,390	39.8	31,337	0.0089
$5,000$	9,693,660	55.3	48,867	1.3
$6,000$	14,594,238	86.7	71,164	1.5
$7,000$	20,494,808	124.3	96,539	1.7
$8,000$	27,395,331	161.2	125,343	1.9
$9,000$	35,295,659	217.8	317,220	2.1
$10,000$	44,195,948	276.9	602,326	2.1

Table 4: Number of Pairs and Runtime comparison on strings and Hamming distance

Figure 5 shows a runtime comparison of MK and MLBA for the case of character strings from a finite alphabet.

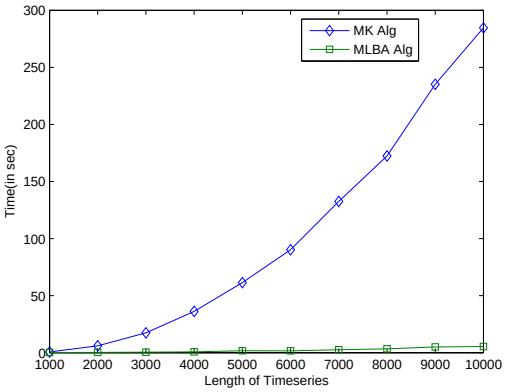

Figure 5: Runtime comparison between MK and MLBA for character strings

5 Identification of the Least Correlated Pair of Strings

The light bulb algorithm of [18] identifies the closest pair of strings, from out of n given binary strings. An interesting question is if we can use the same algorithm to identify the **furthest** pair of strings. This problem has relevance in many problems including the two locus problem in GWAS. The authors of [1] present an elegant adaptation of the light bulb algorithm to solve this problem when the strings are binary. They also show how to solve this problem for arbitrary alphabets using Locality Sensitive Hashing (LSH) [4]. They map the input strings into binary strings using LSH. LSH closely preserves similarities with a high probability. In this section we show how to avoid LSH. In particular, we present novel deterministic mappings of the input strings to binary strings such that similarities are preserved deterministically. Our experimental comparison shows that our algorithm has a significantly better run time than that of [1].

5.1 Some notations

Let $m(x, y)$ stand for the number of matches between two strings (of equal length) x and y. For instance, if $x = 10010$ and $y = 00111$, then $m(x, y) = 2$ (since they match in positions 2 and 4). Let $X = x_1, x_2, \ldots, x_q$ and $Y = y_1, y_2, \ldots, y_q$ be two sequences of strings (each string having the same length). We define $M(X, Y)$ to be $\sum_{i=1}^{q} m(x_i, y_i)$.

Consider the sequences $A_i = a_1^i, a_2^i, \ldots, a_k^i$, for $1 \leq i \leq n$, where each a_j^i is $0, 1$, or 2 (for $1 \leq j \leq k$). Note that each A_i is a sequence of strings where each string is of length 1. Let $M(A_i, A_j) = u_{ij}$.

Each A_i can be thought of as a string from the alphabet $\{0, 1, 2\}$. In the application of GWAS, we can let A_i correspond to the SNP i, for $1 \leq i \leq n$.

Specifically, a_j^i is the value of the ith SNP in subject j, for $1 \leq j \leq k$. If we are interested in finding the two most correlated SNPs, then we can use MLBA to identify this pair (as shown in Section 4). On the other hand, if our goal is to identify the least correlated pair, then, it is not clear how to do this using MLBA. To solve the two locus GWAS problem, we have to identify not only the most correlated pair of bulbs but also the least correlated pair.

5.2 Finding the least correlated pair – the case of zeros and ones

The authors of [1] present an elegant solution for this problem when each A_i has only zeros and ones. In this case, each A_i can be thought of as a light bulb. The idea is to construct a matrix D of size $k \times 2n$ where each column of D corresponds to either a bulb or its 'complement', Specifically, the first n columns correspond to the bulbs and the next n columns correspond to the complements of the bulbs. In other words, $D[j,i] = a_j^i$, for $1 \leq i \leq n, 1 \leq j \leq k$ and $D[j,i] = \bar{a}_j^i$ for $1 \leq j \leq k, (n+1) \leq i \leq 2n$. Here, if x is any bit, then, $\bar{x}$ denotes its complement. Let $D_1 = \{q : 1 \leq q \leq n\}$ and $D_2 = \{q : (n+1) \leq q \leq 2n\}$. The algorithm of [1] for finding the least correlated pair works as follows. Consider all the pairs of columns (a, b) such that $a \in D_1$ and $b \in D_2$. From out of these pairs, identify the pair (a', b') of columns with the maximum number of matches. If $a' = i$ and $b' = n + j$, then (i, j) is the least correlated pair of bulbs. Finding such a pair (a', b') can be done using the light bulb algorithm of [18]. The correctness of this algorithm follows from the fact that if the two bulbs i and j have the least number of matches, then, column i and the complement of column j will have the most number of matches.

5.3 Finding the least correlated pair - the case of zeros, ones, and twos

It is not clear how to extend the above idea when the sequences have three (or more) possible elements. The authors of [1] reduce such general cases to the case of zeros and ones using locality sensitive hashing (LSH). The measure of correlation used by [1] is different from what we use in this paper. We define the correlation between two strings A_i and A_j as $p_{ij} = \frac{M(A_i, A_j)}{k}$. In contrast, [1] use Pearson's correlation coefficient.

In this section we present an elegant algorithm for the problem of identifying the least correlated pair of strings without employing LSH. The idea of [1] is to map input strings into Boolean vectors. If i and j are any two strings, then the sequences A_i and A_j are mapped to Boolean vectors A_i' and B_i' by LSH such that the distance between A_i and A_j will be nearly the same as the distance between A_i' and A_j' with some probability. The larger the length of A_i' is, the better will be the accuracy of LSH in preserving distances.

Our algorithm also maps each A_i into a Boolean vector A_i' deterministically such that $|A_i'| = 3|A_i|$, for $1 \leq i \leq n$.

Consider an alphabet Σ with three strings where $\Sigma = \{001, 010, 100\}$. Clearly, $m(x, y) = 3$ if $x = y$ and $m(x, y) = 1$ if $x \neq y$ for any $x, y \in \Sigma$. Also, $m(x, \bar{y}) = 0$ if $x = y$ and $m(x, \bar{y}) = 2$ if $x \neq y$. Here $\bar{y}$ stands for the string obtained from y by complementing each bit. For example, if $y = 010$ then $\bar{y} = 101$.

Consider the sequences $A_i = a_1^i, a_2^i, \ldots, a_k^i$, for $1 \leq i \leq n$, where each a_j^i is $0, 1$, or 2 (for $1 \leq j \leq k$). Note that each A_i is a sequence of strings where each string is of length 1. Let $M(A_i, A_j) = u_{ij}$. Assume now that we encode each a_j^i as follows (for $1 \leq i \leq n$ and $1 \leq j \leq k$): $0 \rightarrow 001; 1 \rightarrow 010$: and $2 \rightarrow 100$. Let the encoded version of A_i be denoted as A_i' for $1 \leq i \leq n$. Note that $|A_i'| = 3k$, for any $1 \leq i \leq n$.

It is easy to see that $M(A_i', A_j') = 3u_{ij} + (k - u_{ij}) = k + 2u_{ij}$, for any i and j $(1 \leq i, j \leq n)$. For any $A_i' = a_1'^i, a_2'^i, \ldots, a_{3k}'^i$, let $\bar{A}_i' = \bar{a}_1'^i, \bar{a}_2'^i, \ldots, \bar{a}_{3k}'^i$, for $1 \leq i \leq n$. Clearly, $M(A_i', \bar{A}_j') = 2(k - u_{ij})$, for any $1 \leq i, j \leq n$.

Clearly, the following statement is true: If, from out of all the pairs of strings, (i, j) has the largest correlation, i.e., u_{ij} is the largest, then from out of all the Boolean vectors generated, A_i' and A_j' will have the largest correlation. Also, if u_{ij} is the smallest, then, A_i' and $\bar{A}_j'$ will have the largest correlation (from out of the pairs $(A_i', \bar{A}_j'), i \neq j, 1 \leq i, j \leq n$).

We can now form a matrix D of size $3k \times 2n$ where the first n columns correspond to (transformed) strings and the next n columns correspond to complements of (transformed) strings. Let $D_1 = \{q : 1 \leq q \leq n\}$ and $D_2 = \{q : (n + 1) \leq q \leq 2n\}$. Consider all the pairs of columns (a, b) such that $a \in D_1$ and $b \in D_2$. From out of these pairs, identify the pair (a', b') of columns with the maximum number of matches. If $a' = i$ and $b' = n + j$, then (i, j) is the least correlated pair of strings. Finding such a pair (a', b') can be done using the light bulb algorithm of [18].

5.4 Run Time Analysis

Theorem 5.1. *Given n strings, we can find the closest pair of strings in an expected time of $O\left(n^{1+\frac{\log p_1}{\log p_2}} \log^2 n\right)$, where p_1 and p_2 are the largest and the second largest correlation values, respectively. Also, we can find the least correlated pair of strings in an expected time of $O\left(n^{1+\frac{\log((2/3)(1-c_1))}{\log((2/3)(1-c_2))}} \log^2 n\right)$, where c_1 and c_2 are the smallest and the next smallest correlation values, respectively.*

Proof: When we transform input strings to binary sequences, the ordering of pairs is preserved in terms of correlations as we have shown before. Let p_1 be the correlation of the largest correlated pair and p_2 be the correlation of the second largest correlated pair. How do these values change in the transformed domain? If p_1' and p_2' are the transformed values of these correlations, respectively, it can be seen that $p_1' = \frac{1}{3} + \frac{2}{3}p_1$ and $p_2' = \frac{1}{3} + \frac{2}{3}p_2$.

If c_1 and c_2 are the correlations of the smallest and the second smallest correlated pairs, respectively, and if c_1' and c_2' are the transformed values of these, respectively, then we can see that: $c_1' = \frac{2}{3}(1 - c_1)$ and $c_2' = \frac{2}{3}(1 - c_2)$. To

find the largest correlated pair, we can use MLBA (Theorem 4.1). We use the mapping only to find the least correlated pair. □

5.5 The case of a general alphabet

We have thus far considered the case where the alphabet is $\{0, 1, 2\}$. We can extend the mapping to a general alphabet and get the following theorem.

Theorem 5.2. *Given n strings, we can find the largest correlated pair of strings in an expected time of $O\left(n^{1+\frac{\log p_1}{\log p_2}} \log^2 n\right)$, where p_1 and p_2 are the largest and the second largest correlation values, respectively. Also, we can find the least correlated pair of strings in an expected time of $O\left(n^{1+\frac{\log((2/\sigma)(1-c_1))}{\log((2/\sigma)(1-c_2))}} \log^2 n\right)$, where c_1 and c_2 are the smallest and the next smallest correlation values, respectively, and there are σ characters in the alphabet.*

Proof: Consider sequences from the alphabet $\{0, 1, \ldots, \sigma - 1\}$. In this case we map each element of this alphabet to a binary string of length σ where there is only one 1. Specifically, we use the following mapping: $0 \rightarrow 00\cdots001$; $1 \rightarrow 00\cdots010$; etc. As before, we don't need any mapping if our goal is to find the largest correlated pair. The mapping is used only to find the least correlated pair. □

We can improve the above theorem by employing a random mapping as follows: We will use a binary string of length σ to encode each symbol in the alphabet. The encoding for each symbol is obtained by (uniformly) randomly choosing each bit in the string (of length σ). Let x and y be any two symbols in the alphabet (with $x \neq y$) and let e_x and e_y be their encodings, respectively. Then, clearly, the expected value of $m(e_x, e_y)$ is $\frac{\sigma}{2}$. Also, the expected value of $m(e_x, \bar{e}_y)$ is $\frac{\sigma}{2}$. If c is the correlation between a pair of strings and if c' is the transformed value, then, it follows that the expected value of c' is $\frac{1}{2}(1 - c)$. An application of the Chernoff bounds will readily imply that the value of c' will indeed be very close to this expected value with a probability of $1 - \sigma^{-\Omega(1)}$. Therefore, we get:

Theorem 5.3. *Given n strings, we can find the least correlated pair of strings in an expected time of $O\left(n^{1+\frac{\log((1/2)(1-c_1))}{\log((1/2)(1-c_2))}} \log^2 n\right)$, where c_1 and c_2 are the smallest and the next smallest correlation values, respectively, and there are σ characters in the alphabet.* □

6 Two Locus Association Problem

The two locus association problem is defined as follows. Input is a matrix M of size $(m_1 + m_2) \times n$ where $m_1 + m_2$ is the number of patients (subjects) each with n SNPs. Here m_1 is the number of cases and m_2 is the number of controls. There are three possible values for each SNP, namely, $0, 1$, or 2. The cases are of phenotype 1 and the controls are of phenotype 0. Rows 1 through m_1 of M correspond to cases. Let this submatrix be called A. Rows $m_1 + 1$ through m_1+m_2 of M correspond to controls and let this submatrix be called B. Each column of M corresponds to an SNP. The two locus association problem is to identify the pair of SNPs whose statistical correlation with phenotype is maximally different between cases and controls. As mentioned in [1], the goal is to identify the pair:

$$\arg\max_{i,j} |P_A(i,j) - P_B(i,j)|.$$

If Q is any matrix, then, $P_Q(i,j)$ stands for the correlation between the columns i and j of Q.

The algorithm of [1] exploits the light bulb algorithm of [18] and locality sensitive hashing (LSH) [4]. They use LSH to transform matrices A and B to A' and B', respectively. In particular, each column c_i of A is converted to a column c'_i of zeros and ones. The size of c_i is $1 \times m_1$ and the size of c'_i is chosen to be $u = \max\{m_1, m_2\}$. The matrix B is also transformed into B' in a similar manner using LSH. Followed by this, the pair of interest is identified.

To be precise, using A' and B', the matrix D is formed where

$$D = \begin{bmatrix} A' & A' \\ B' & \bar{B}' \end{bmatrix}$$

where $\bar{B}'$ is obtained from B' by complementing every element of B'. Note that D is of size $2u \times 2n$. Let $D_1 = \{1, 2, \ldots, n\}$ and $D_2 = \{n+1, n+2, \ldots, 2n\}$. Consider all the pairs of columns (i, j) such that $i \in D_1$ and $j \in D_2$. From out of these pairs, identify the pair (i', j') of columns with the maximum number of matches. If $i' = a$ and $j' = n + b$, then (a, b) is the pair of interest. This pair can be found using the light bulb algorithm of [18].

We can use our mapping ideas to get the following theorem.

Theorem 6.1. *We can find the pair (i, j) of SNPs that maximizes $P_A(i, j) - P_B(i, j)$ in an expected time of $O\left(n^{1+\frac{\log((1/2)+(p_1/3))}{\log((1/2)+(p_2/3))}} \log^2 n\right)$, where p_1 and p_2 are the smallest and the next smallest values of $P_A(i, j) - P_B(i, j)$, respectively, over all possible pairs (i, j) of SNPs.*

Proof: Our algorithm also uses the same method except that instead of using LSH to map A and B to A' and B', respectively, we employ the deterministic mapping we have proposed in Section 5.3.

Let i and j be two SNPs (i.e., two columns in A and B). Let $P_A(i, j) = p$ and $P_B(i, j) = q$. Now consider columns i and $n + j$ of D. What can we say

about the correlation of these columns? From the discussion in Section 5.4, we realize that this correlation is $\frac{1}{2} + \frac{1}{3}(p - q)$. This also proves the correctness of our algorithm. Let p_1 be the maximum value of $P_A(i, j) - P_B(i, j)$ over all possible pairs (i, j) and let p_2 be the second largest value. Then, the run time follows from Theorem 4.1. □

In a similar manner we can also find the pair that maximizes $P_B(i, j) - P_A(i, j)$ and hence identify the pair that maximizes $|P_A(i, j) - P_B(i, j)|$.

6.1 An experimental comparison

The notion of similarity (between two SNPs) used in [1] is Pearson's correlation coefficient. In this paper the similarity we use is based on the Hamming distance. Specifically, the complement of the Hamming distance is the measure of similarity we employ. The authors of [1] have tested their algorithms on different data sets (including random data). Since we do not have access to either these data sets or their algorithms, the only comparison we can do was on the random data. As explained in [1], we have also generated SNPs from binomial distributions. In particular, for each subject, the value of each SNP is chosen uniformly randomly to be either 0 or 1 with equal probability. This dataset is called NOISE Data in [1] (c.f. Table 6 in [1]). Like in [1], we have also generated data of sizes 10K, 50K, and 100K. For each size we have generated two different data sets and computed the average number of pairs to be processed. We compare these numbers with the ones reported in [1]. As can be seen from Table 5, our algorithm is around 4 times faster than the one in [1]. Note that this a significant improvement since the typical processing times for the two locus problem are quite high. For example, the authors of [1] report that on some of the data sets (with no more than 10^7 SNPs), the brute force algorithm for the two locus problem took several days on 1000 CPUs! Thus any improvement in the run time could make a noticeable difference.

How does one ensure that the output of an algorithm for the two locus problem is correct? For small data sizes, one could run the exhaustive brute force algorithm to identify the correct pair and use it to verify correctness. In fact when the number of SNPs is either 10K or 50K, we first found the correct answer and then used it to measure the run time of our algorithm as follows. We'll run our algorithm one iteration at a time until the correct pair(s) is (are) picked up by our algorithm. At this point we will stop and report the total number of pairs collected. The numbers shown in Table 5 have been obtained in this manner. We could not use this method for $n >$ 50K, since the brute force algorithm was taking too much time.

When n is very large, we *inject* pairs with known correlations. As an example, consider the problem of finding the largest correlated pair of columns in a $m \times n$ matrix A. Say we generate each column by picking each element to be either 0 or 1 with equal probability. For any two columns, clearly, the expected correlation is $\frac{1}{2}$. We can perform a probabilistic analysis to get a high probability bound on the largest correlation between any two columns. One could also get this estimate empirically. For example, for $n = 10,000$, we generated several

random data sets and computed the largest correlation in each and calculated an average. The average maximum correlation was 58.3%. Let p be this value. To inject a pair with a correlation of p' where p' is $> p$ we generate a column a with all ones and another column b with $p'm$ ones and $m(1-p')$ zeros. We then replace (any) two columns of A with a and b. Clearly, the correlation between a and b is p'. The expected correlation between a and any other column of A (other than b) is $\frac{1}{2}$. Similarly, the expected correlation between b any other column of A (other than a) is $\frac{1}{2}$. Thus the pair (a, b) is likely to be the winner with high probability. We stop our algorithm when this pair is picked by our algorithm. We have picked a value for p' that is only slightly larger than p so as to get an accurate estimate on the run time.

We have used a similar technique to inject pairs for the two locus problem as well. In Table 6 we show the results for our algorithm. In these cases we have generated the SNPs randomly from a binomial distribution as before. We have employed 200 cases and 200 controls. Clearly, the expected value of $(p-q)$ is zero for this data. If we map this data using our deterministic mapping, then the expected correlation between any two SNPs will be $\frac{1}{2}$. Here again we empirically found that the largest correlation was around 58.3% (when n was 10K). Therefore, we have injected a pair whose correlation was 60%. This pair is likely to be the winner. Our algorithm was run until this pair was picked. At this time, the algorithm was stopped. We also checked if the algorithm picked any pair whose correlation was better than that of the injected one and found none. In all the datasets we tried, we were always able to find a pair (other than the injected one) whose correlation was very close to 60% and hence the numbers shown in Table 6 are very close to the case with no injections.

Dataset Size	Pairs (GWAS)	Pairs (MLBA)	Top 10	Top 100
$10,000$	7,082,458	1,904,999	0.6	0.53
$50,000$	13,626,181	53,907,259	0.4	0.34

Table 5: No of pairs comparison on NOISE data set

Dataset Size	No. of Pairs (MLBA)
$10,000$	560426
$50,000$	706197

Table 6: No of pairs comparison on synthetic data with injected pairs

Practical Considerations: Another important question is for how long we

should run the program before we can be sure that the correct pair has been obtained (with a high confidence). Please note that we do not know the values of p_1 and p_2 (c.f. Theorem 6.1). Theorem 6.1 suggests that the run time of our algorithm is $O(n^{1+\alpha})$ for some relevant α. We can empirically estimate α. The idea is to measure the run time of the algorithm for various values of n (as explained above), the maximum value of n being as much as possible (for the given computing platform and time constraints). Then we could use any numerical procedure to estimate α.

7 Conclusions

In this paper we have presented novel algorithms for the closest pair problem (CPP). CPP is a ubiquitous problem that has numerous applications in varied domains. We have offered algorithms for the cases of Euclidean as well as Hamming distances. We have applied our algorithms for two well studied and important problems, namely, time series motif mining and two locus genome wide association study. Our algorithms significantly improve the best-known algorithms for these problems. Specifically, we improve the results presented in many prior papers including [15], [5], and [1].

8 Acknowledgement

This research has been supported in part by the NIH grant NIH-R01-LM-010101.

References

[1] P. Achlioptas, B. Scholkopf and K. Borgwardt, Two-locus association mapping in subquadratic runtime, ACM SIGKDD international conference on Knowledge discovery and data mining (KDD) 2011.

[2] C.E. Aston, D.A. Ralph, D.P. Lalo, S. Manjeshwar, B.A. Gramling, D.C. DeFreese, A.D. West, D.E. Branam, L.F. Thompson, M.A. Craft, D.S. Mitchell, C.D. Shimasaki, J.J. Mulvihill, and E.R. Jupe, Oligogenic combinations associated with breast cancer risk in women under 53 years of age, *Human Genetics*, 116(3):208-221, Feb. 2005.

[3] P. Beaudoin, M. van de Panne, P. Poulin and S. Coros, Motion-Motif Graphs, Symposium on Computer Animation 2008.

[4] M. S. Charikar, Similarity estimation techniques from rounding algorithms, Proc. ACM Symposium on Theory of Computing (STOC), 2002.

[5] B. Chiu, E. Keogh, and S. Lonardi, Probabilistic discovery of time series motifs, Proc. of the 9th International Conference on Knowledge Discovery and Data mining (KDD), pp. 493-498, 2003.

[6] J.H. Cho, D.L. Nicolae, L.H. Gold, C.T. Fields, M.C. LaBuda, P.M. Rohal, M.R. Pickles, L. Qin, Y. Fu, J S. Mann, B.S. Kirschner, E.W. Jabs, J. Weber, S.B. Hanauer, T.M. Bayless, and S.R. Brant. Identification of novel susceptibility loci for inflammatory bowel disease on chromosomes 1p, 3q, and 4q: evidence for epistasis between 1p and IBD1 Proceedings of the National Academy of Sciences of the United States of America, 95(13):7502-7507, June 1998.

[7] T. H. Consortium. A second generation human haplotype map of over 3.1 million SNPs. Nature, 449(7164):851-61, Oct. 2007.

[8] H. J. Cordell. Detecting gene-gene interactions that underlie human diseases. Nat Rev Genet, 10(6):392-404, June 2009.

[9] N.J. Cox, M. Frigge, D.L. Nicolae, P. Concannon, C.L. Hanis, G.I. Bell, and A. Kong, Loci on chromosomes 2 (NIDDM1) and 15 interact to increase susceptibility to diabetes in mexican americans, Nature Genetics, 21(2):213-215, Feb. 1999.

[10] S. Fortune and J. S. Hopcroft, *A note on Rabin's nearest-neighbor algorithm*, Information Processing Letters, 8(1), pp. 20-23, 1979.

[11] T. Guyet, C. Garbay and M. Dojat, Knowledge construction from time series data using a collaborative exploration system, Journal of Biomedical Informatics 40(6): 672-687 (2007).

[12] S. Khuller and Y. Matias, *A Simple Randomized Sieve Algorithm for the Closest-Pair Problem*, Information and Computation, Vol 188 (1), pp. 34–37, 1995.

[13] J. Meng, J.Yuan, M. Hans and Y. Wu, Mining Motifs from Human Motion, Proc. of EUROGRAPHICS, 2008.

[14] D. Minnen, C.L. Isbell, I. Essa, and T. Starner, Discovering Multivariate Motifs using Subsequence Density Estimation and Greedy Mixture Learning, Proc. 22nd Conference on Artificial Intelligence (AAAI?07), 2007.

[15] A. Mueen, E. Keogh, Q. Zhu, S. Cash, B. Westover, *"Exact Discovery of Time Series Motifs"*, SIAM International Conference on Data Mining (SDM) 2009.

[16] S. K. Musani, D. Shriner, N. Liu, R. Feng, C. S. Coffey, N. Yi, H. K. Tiwari, and D. B. Allison. Detection of gene x gene interactions in genome-wide association studies of human population data. Human Heredity , 63(2):67-84, 2007.

[17] R. Nakamichi, Y. Ukai, and H. Kishino, Detection of closely linked multiple quantitative trait loci using a genetic algorithm, Genetics, 158(1):463-475, May 2001.

[18] R. Paturi, S. Rajasekaran, and J. Reif. The light bulb problem. In Proc. 2nd Annu. Workshop on Comput. Learning Theory , pp. 261-268, San Mateo, CA, 1989. Morgan Kaufmann.

[19] F. Preparata and M. Shamos, *Computational Geometry*, Springer Verlag, 1986.

[20] M. Rabin, *Probabilistic Algorithms*, Algorithms and Complexity, Recent Results and New Directions, Academic Press, pp 21-39, 1976.

[21] S. Rombo and G. Terracina, Discovering representative models in large time series databases, Proc. 6th International Conference on Flexible Query Answering Systems, pp. 84-97, 2004.

[22] Y. Tanaka, K. Iwamoto, and K. Uehara, Discovery of time-series motif from multi-dimensional data based on MDL principle, Machine Learning, 58(2-3):269-300, 2005.

[23] S. Tata, Declarative Querying For Biological Sequences, Ph.d Thesis, The University of Michigan, 2007. (Advisor Jignesh M. Patel).

[24] Y. Wang, X. Liu, K. Robbins, and R. Rekaya, AntEpiSeeker: detecting epistatic interactions for case-control studies using a two-stage ant colony optimization algorithm, BMC Bioinformatics, 3:117-117, 2010.

[25] J. Xu, C.D. Langefeld, S.L. Zheng, E.M. Gillanders, B. Chang, S.D. Isaacs, A.H. Williams, K E. Wiley, L. Dimitrov, D.A. Meyers, P.C. Walsh, J.M. Trent, and W.B. Isaacs, Interaction effect of PTEN and CDKN1B chromosomal regions on prostate cancer linkage, Human Genetics, 115(3):255-262, Aug. 2004.

[26] A.C. Yao, *Lower Bounds for algebric computation trees with integer inputs*, SIAM J. Comput., 20:4, pp. 655-668, 1991.

[27] X. Zhang, S. Huang, F. Zou, and W. Wang. TEAM: efficient two-locus epistasis tests in human genome-wide association study. Bioinformatics (Oxford, England), 26(12):i217-227, June 2010.

[28] X. Zhang, F. Zou, and W. Wang, Fastanova: an efficient algorithm for genome-wide association study, in Proc. 14th ACM SIGKDD international conference on Knowledge discovery and data mining (KDD), pages 821-829, Las Vegas, Nevada, USA, 2008.